



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-SR-411009

Final Report for Subcontract B541028, Pore-Scale Modeling to Support "Pore Connectivity" Research Work

R. P. Ewing

March 3, 2009

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Final Report for Subcontract **B541028**,
Pore-Scale Modeling to Support "Pore Connectivity" Research Work
PI: Robert P. Ewing, Iowa State University

For the project

Pore Connectivity, Episodic Flow, and Unsaturated Diffusion in Fractured Tuff
PI: Qinhong “Max” Hu, Lawrence Livermore National Laboratory

Introductory Comments

This report covers modeling aspects of a combined experimental and modeling task in support of the DOE Science and Technology Program (formerly OSTI) within the Office of Civilian Radioactive Waste Management (OCRWM).

Research Objectives

The research for this project dealt with diffusive retardation: solute moving through a fracture diffuses into and out of the rock matrix. This diffusive exchange retards overall solute movement, and retardation both dilutes waste being released, and allows additional decay.

Diffusive retardation involves not only fracture conductivity and matrix diffusion, but also other issues and processes: contaminants may sorb to the rock matrix, fracture flow may be episodic, a given fracture may or may not flow depending on the volume of flow and the fracture’s connection to the overall fracture network, the matrix imbibes water during flow episodes and dries between episodes, and so on.

The objective of the project was to improve understanding of diffusive retardation of radionuclides due to fracture / matrix interactions. Results from combined experimental / modeling work were to (1) determine whether the current understanding and model representation of matrix diffusion is valid, (2) provide insights into the upscaling of laboratory-scale diffusion experiments, and (3) help in evaluating the impact on diffusive retardation of episodic fracture flow and pore connectivity in Yucca Mountain tuffs.

Questions explored included the following:

- What is the relationship between the diffusion coefficient measured at one scale, to that measured or observed at a different scale? In classical materials this relationship is trivial; in low-connectivity materials it is not.
- Is the measured diffusivity insensitive to the shape of the sample? Again, in classical materials there should be no sample shape effect.
- Does sorption affect diffusive exchange in low-connectivity media differently than in classical media?
- What is the effect of matrix saturation on the effective diffusivity? Is it different for low-connectivity media than for classical media?
- In addition to changing the matrix saturation (and thereby the diffusion coefficient), the wetting/drying cycles drive water into, and then out of, the matrix. How do these mass flow cycles affect the long-term exchange of solutes between the fracture and the matrix? Can it

be treated as a simple increase in effective diffusivity? Is it a local or a global effect? Is the effect different in low-connectivity media?

The modeling portion of this project primarily focused on how diffusion varies with pore connectivity, and it also connected the experimental work to theory.

Approach

Experimental Work

Various properties of YMP rock were measured; some other rocks were also examined for comparison. Several procedures were used and compared:

- (1) Dry rock cores, surface-epoxied except at the top and bottom, were suspended from an electronic balance over a water reservoir. The mass of water imbibed up into the core was recorded as a function of time, starting when the core first touched the water surface. Effects of buoyancy, evaporation, and drift were eliminated. This, the most basic (and in some ways the most revealing) of all the experiments (see accomplishments, below), was performed for different size and shape samples of several rocks.
- (2) The measurements in (1) were also performed using a tracer rather than pure water. The tracer concentration as a function of distance from the imbibing face was then measured using LA/ICP-MS. This was done with both sorbing and non-sorbing tracers.
- (3) Tracer movement was evaluated over the course of four wetting and drying cycles, with a different suite of both sorbing and non-sorbing tracers added at the beginning of each wetting cycle. The final disposition of the tracers was evaluated at the end of all four cycles, using LA/ICP-MS. One experiment with 4 episodes was completed.

Two other experimental portions of the overall project were the following:

- (4) Synchrotron microtomography of a KI tracer, allowing analysis of the diffusion front. A custom sample stand was built, and some data collected, but the spatial resolution of the diffusion front was not high enough to meet the objectives.
- (5) Independent gas-phase measurement of the diffusion coefficient of the rock samples. The method to be used involved diffusive counter-flow of two gases, but a persistent gas pressure gradient across the sample apparently prevented the measurements from giving a pure diffusion coefficient.

Modeling Work

Because a significant role of the modeling was to connect the experimental results to theory, the relevant theory should be identified. The statistics of low-connectivity media is the province of **percolation theory** (Stauffer and Aharony, 1994; Hunt, 2005). It is not the purpose of this report to present a primer on this theory, but some basic concepts are introduced in Appendix 1.

Many of the experiments were modeled using a pore-scale network model. The development of the model used is documented in the project proposal, statement of work, and scientific notebook (SN); source code is also provided on a supplementary CD. Briefly, the original model (Ewing and Berkowitz, 2001) was an internal diffusion-limited aggregation (also called anti-DLA and diffusion-limited annihilation; Meakin and Deutsch, 1986) growth model: a simple cubic lattice porespace was generated and stochastically pruned to the desired connectivity, and all pores were considered “empty” except those at the inlet face. Random walkers were successively released from the inlet face; when a walker first encountered an empty pore it stopped, filling it. The procedure is a simple but extremely CPU-intensive analog of imbibition: it might easily require over a million random walkers before the wetting front reached the far face, and successive walkers require increasingly long times (often billions of time steps) to find an empty pore.

A second-generation model was therefore developed to speed the simulations. This model released a fixed number of random walkers (usually 1000) at the inlet face, and recorded each particle’s first arrival time at every (discrete) distance from the inlet face. Results were statistically identical with the original model (see file *testimbibe2.xls*), so this second generation model was used for most of the simulations. Unfortunately, neither model directly addressed issues of matrix saturation, as discussed below.

Experiments and simulations both produced results that, when plotted in logarithmic space (log cumulative mass imbibed as a function of log time), produced one of three results: (1) a straight line with a slope of 0.5 (because imbibition and diffusion are classically proportional to the square root of time; see Philip, 1957, and Bruce & Klute, 1956), (2) a slope of approximately 0.265, or (3) a slope of about 0.265, transitioning to a slope of 0.5. A persistent issue in the analysis was rigorously identifying the breakpoint between the two slopes. The experimenter had a relatively small number of such cases, and was content to fit them by eye. In order to run Monte Carlo simulations with (generally) 100 realizations and many “pore connectivity” values, it was desirable to use a method that was objective but not time-intensive. Satisfactory procedures were identified and tested for this purpose. The programs written during the development of these procedures have been provided to LLNL in addition to this report. A brief summary of breakpoint statistics is given in Appendix 2.

Results

The main result of the modeling was that it successfully connected the experimental work to existing theory. Several features in the experimental data, not explained by more classical theories, were both evident in the simulation output and consistent with predictions of percolation theory. Below, there is first a description of how the simulations are related to the experiments, followed by a discussion of several of the data features and conceptual connections.

A geological porous medium has pores, some or all of which are active, that is, connected to other pores. Call the intersections pore bodies, and the connections pore throats. The parallels between porous media, network models, and graph theory are simple: pore bodies correspond to the sites (intersections) in a graph or a pore network model, and pore throats are the bonds. Where an active bond exists between neighboring sites, it denotes a pore throat connecting two pore bodies. If (for example) a simple cubic lattice had every site connected by active bonds to

its nearest neighbors, it would represent a medium in which all pores were accessible; this medium would have a mean pore coordination of 6.

Two conditions must be met for a solute to diffuse from one pore to a neighboring pore: (1) there must be an open pore throat (an active bond) connecting the two pores, and (2) both the pores and the pore throats must be filled with water. In other words, the presence (absence) of water plays the same role as the presence of an active (missing) pore throat. Consequently, pore connectivity and matrix saturation can be equivalently adjusted by the single model parameter p , the connection probability. This conceptual equivalence can be seen in e.g. Hunt (2005), where the probability p in site or bond percolation may be replaced by either the volume water content θ or the porosity ϕ in continuum percolation, as circumstances warrant.

Overview: comparison of imbibition simulation with imbibition experiment

An imbibition experiment produces data showing a curve with up to 4 regions (Fig. 1). These are referred to as the initial sorption region, the low-slope region, the high-slope region, and the final plateau region. The initial sorption region shows rapid initial uptake of water during the first few seconds of the experiment, with water having essentially no viscous resistance. That is, the water is filling pores and sorbing to hydrophilic surfaces that are exposed on the face: no travel through the medium is required. This initial phase is rapid (perhaps 20 s in Fig. 1), and involves very little water. The simulations produce equivalent patterns (Fig. 2), though of course the specific masses and times differ.




Figure 1. Cumulative imbibition into tuff sample TSw34L8, showing initial sorption, low-slope, and high-slope regions. Max gives low slope = 0.27, and high slope = 0.42. This sample was not run until the wetting front reached the top: there is no plateau.

The low-slope region, which is the first segment if both low- and high-slope regions are present, corresponds to imbibition before the wetting front has penetrated one correlation length χ from the inlet face (see Appendix 1). At distances $l < \chi$, water coming from the inlet face encounters finite clusters that are accessible only because the medium has a cut face. This gives a high accessible porosity that decreases as the wetting front advances; the tortuosity encountered by the wetting front also increases with distance. The medium has fractal character at this point, and the imbibition slope takes a theoretical slope of approximately 0.265 (Stauffer and Aharony, 1994).

For distances $l > \chi$, only the infinite cluster is accessible, and with χ established as the effective length scale, the medium now behaves classically. That is, the imbibition slope goes to 0.5, and the tortuosity holds to a constant value. This is the high-slope region, also called the second

segment when both low- and high-slope regions are present¹. The finite clusters that aren't accessible from the inlet face are now “holes” in the infinite cluster. As Stauffer and Aharony (1994) picturesquely put it, at scales “larger than the typical hole size [χ], the [diffusing molecule] feels only an average over the small holes, just as the tyres of your Rolls-Royce average over the small pores of the asphalt over which your chauffeur is driving you.”




Figure 2. Simulated imbibition with connectivity $p = 0.25$ (one realization).

As imbibing water reaches the top of a core, there are fewer and fewer empty pores to fill, so the imbibition levels off. This appears as a plateau in the imbibition curve. Again, the same behavior shows up in the simulations. These simulations have not been run in a more “mainstream” simulation model, such as Hydrus, to see if it reproduces any of this behavior. It may be expected that a mainstream model would reproduce the high-slope region and the end plateau, and it might reproduce the initial sorption region, but it certainly wouldn't reproduce the low-slope region.

In summary, the general pattern is that, at high connectivities, the imbibition slope of 0.5 was seen. As p decreased toward the percolation threshold p_c , the slope also decreased. Quite close to p_c , the observed slope was either the low slope (about 0.265) throughout, or it started low and then transitioned to the classical 0.5. For $p < p_c$, imbibition would proceed only a short distance before all pathways were blocked. It is emphasized here that the low-slope region is not “programmed into” the model: there is nothing in the code that examines the pore connectivity and “decides” what slope to produce. The behavior of the imbibition slope is an emergent property of random walks on a low-connectivity network.

¹ The initial sorption region is not considered to be a “segment”: it is basically a noisy nuisance. All data for times < 1000 time steps was generally discarded. Additionally, the second generation simulations don't have a final plateau, and the first generation simulations are often halted by then. So there may be up to four regions, but in terms of finding a breakpoint, there are only one or two segments.

Matrix saturation

Several simulations were run using the first generation (anti-DLA) program, to assess whether the program saw any difference (other than the actual imbibition rate) in imbibition starting from different antecedent (initial) water contents. These anti-DLA simulations require huge amounts of CPU time, so only a few were run, and generally on fairly small (e.g., 64^3) networks. Starting imbibition from different water contents did not affect the information gathered about the medium. When starting with a moister medium, the program took less time but the data were not as clean. Experiment and simulation regarding changes in imbibition rate as a function of initial saturation were not directly compared.

Imbibition slopes and changes in slope

The imbibition slopes that will theoretically be encountered are 0.265 and 0.50. Neither experiment nor simulation was strictly confined to this choice; the range of values found is attributable to finite-size effects. On average, however, the slopes adhered closely to the theoretical values (Fig. 3). Broadly speaking, a wider range of slopes was seen in smaller networks, and a narrower range in experiments than in simulations. This latter observation reflects the samples' size – thousands or millions of pores across, in contrast to the several hundred “pores” across a network model.




Figure 3. Mass imbibed as a function of time, using the second generation program. Connectivity decreases from $p = 1.0$ (left-most line) to $p = p_c$ (right-most line). Log gridlines show that slopes are generally either 0.5, or approximately 0.265.

The greatest range and variability in slope was seen in high-slope regions that followed a low-slope region. Here again the explanation is finite-size effects: when a low-slope region is present, the relevant scale of the system is the correlation length χ rather than the mean pore size (see Appendix 1 for some percolation concepts). The closer a system is to the percolation

threshold, the smaller the sample is *in units of χ* . The variability of that portion of the system encountered beyond χ is therefore quite high, as shown in Figure 4.




Figure 4. Standard deviation of the mean first arrival time of 5000 random walkers at 256 lattice units from the inlet face. The trend (drawn by eye) changes around $p - p_c = 0.01$ ($p = 0.26$), below which value a non-negligible low-slope region is first encountered.

It would be convenient to be able to estimate the distance to the changepoint, χ , by assuming that it is that fraction of the total sample length given by the fraction of water imbibed at the change in slope. Unfortunately, that would be incorrect: the accessible porosity ϕ_a is greatest at the inlet face, and decreases to a constant value for distances $L > \chi$ (see Appendix 1). But fortunately, the accessible porosity ϕ_a at a given distance l from the inlet face *is* related to distance by the

(theoretical) relationships $\phi_a(l) = \phi_a|_{l=0} l^{-\beta/\nu}$ for $l < \chi$, and $\phi_a(l) = \phi_a|_{l=0} \chi^{-\beta/\nu}$ for $l \geq \chi$. In these equations, β and ν are exponents from percolation theory; in 3D, $\beta/\nu \approx 0.466$. This makes estimating χ from the imbibition data slightly more cumbersome, but it is doable.

Sample size

There is little point to running second generation simulations on small lattices, for two reasons. First, as mentioned above, the variability is at least partially due to finite size effects. Second, if looking for behaviors of large ($L > \chi$) systems, one needs to focus on systems for which L is large relative to χ , and yet χ is large enough to have some effect. Some otherwise identical simulations were run on lattices of size 64^3 and 256^3 . The decrease in slope with decreasing connectivity is similar across lattice size when connectivities aren't too low (Fig. 5). However,

differences do show up at very low connectivities, which (inconveniently!) constitute the region of interest for this project. This illustrates the potential for finite-size effects, again highlighting the importance of using large lattices; related issues are discussed below under “shape”.

Sample shape

Because earlier work by Ewing and Gupta (1993) and Ewing and Horton (unpublished) had indicated that there might be a sample shape effect for media near the percolation threshold, it was viewed as useful to examine this possibility in more detail. One simple way to quantify the sample shape is with the aspect ratio, the ratio of the height to the diameter; for the simulations it is the x -dimension divided by the y - or z -dimensions (which are identical). However, because computer resources are limited, many simulations were run using an approximately constant total volume, and simply varying the shape. The resulting data confound the size and shape effects, as seen in Figure 6 (identical to Figure 5, but with other aspect ratio simulations added in). The lattice for the high aspect ratio simulation (height/width = 4) is 512 lattice units tall; the aspect ratio 0.25 simulation used a lattice 102 units tall. In such cases, it is difficult to distinguish between the effects of size and shape.




Figure 5. Difference between fitted first-segment slope and 0.265 (the *theoretical* low or first segment slope) as a function of connectivity $p - p_c$ for two different size lattices. At high p values the slope is 0.5, so points are at $0.5 - 0.265 = 0.235$.




Figure 6. Mean slope of single-segment imbibition curves as a function of connectivity $p - p_c$, for several different aspect ratios and volumes. Data from Figure 5 (ratio = 1) are included for comparison.

When one examines only those simulations performed at $p = p_c$, shape effects are easier to isolate. Additionally, these graphs (Figure 7) use only data from simulations with $x = 128$. Samples with width = height more generally have a low initial slope, but only rarely have two distinct slopes. Thin samples have lower *apparent* correlation length. The correlation length is a

property of the material, not of the sample's shape. What is happening is that the apparent χ is constrained by the system size: in a narrow system, it cannot exceed the system width. But in fact, at $p = p_c$ a tall thin system should not percolate at all. The high aspect ratio simulations shown in Figure 7 form a biased sample, comprised of only those unlikely lattices that *did* percolate. (For example, at $p = p_c$ and an aspect ratio of 4 (512 high, 128 wide in y and z), only 12 in 1000 simulations would percolate.) The trends in Fig. 7 cannot therefore be taken as representative. Samples that did not percolate would have shown a low slope, followed by a plateau – no increase in mass or distance – long before the wetting front reached the top of the core.

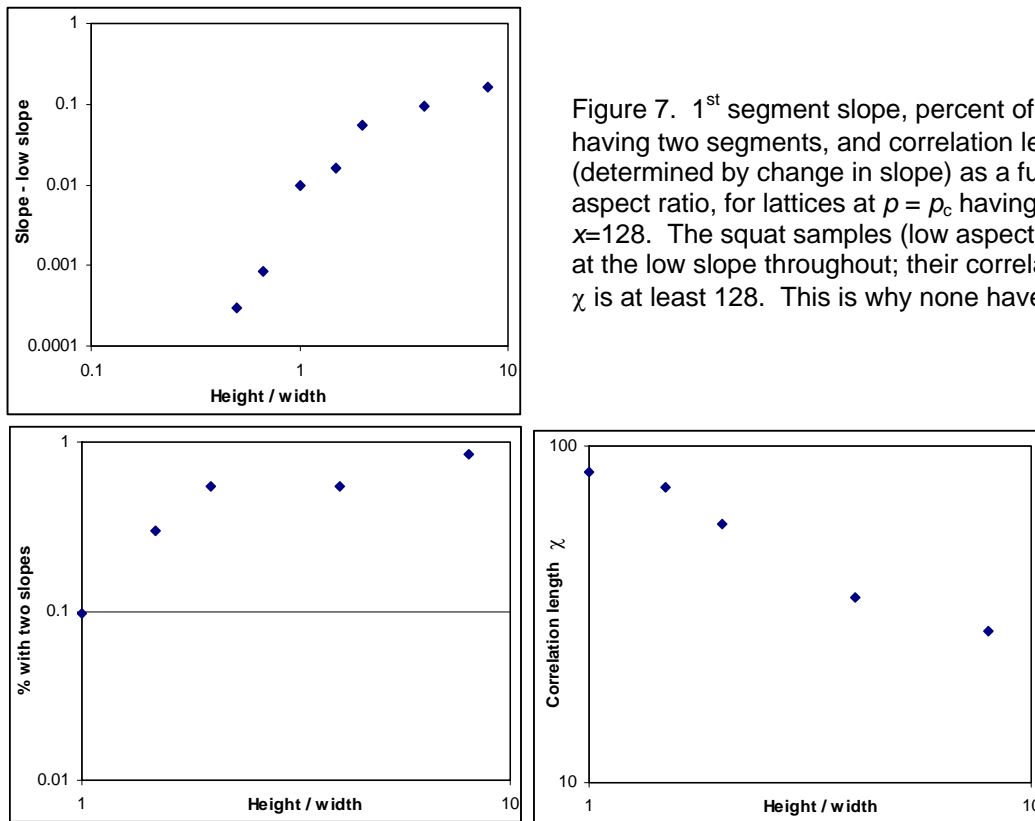


Figure 7. 1st segment slope, percent of simulations having two segments, and correlation length χ (determined by change in slope) as a function of aspect ratio, for lattices at $p = p_c$ having height $x=128$. The squat samples (low aspect ratio) stay at the low slope throughout; their correlation length χ is at least 128. This is why none have two slopes.

Suppose one takes a moderately well-connected sample: it percolates and has a high imbibition slope. Repeatedly core the sample to half its previous radius, dry it, epoxy it along the sides, and conduct an imbibition test. Eventually the sample will no longer percolate (Figure 8). The last sample to percolate may show two distinct slopes, but it also may not: among the hundreds of simulations run, only a small percentage showed two distinct slopes. But even without those two slopes, something about the correlation length can be learned. Because the correlation length is determined by the connectivity of the material rather than by the shape of the sample, one can use the width of the sample to determine χ . Specifically, the correlation length for a given material is bracketed by the width of the last sample to percolate, and the first sample to not percolate. This illustrates why the correlation length is sometimes also called the “mean path separation”: that a cross-sectional area smaller than χ^2 would be unlikely to contain a pathway on

the infinite cluster. Notice that, for the physical experiments, this method requires that the sample properties not be disturbed by the coring (e.g., by cracking), and that each smaller core remain representative of the material as a whole (i.e., that the material be homogeneous at the scale of the sample).




Figure 8. *Left Panel:* Distribution of the maximum accessible distance in a $x=512$ lattice at $p = p_c$, for different widths (y, z). At $x=y=z$, approximately 50% percolate. The median (50th %ile) distance is approximately equal to the lattice width, over two orders of magnitude. *Right Panel:* Percent percolating as a function of width for several connection probabilities. Repeatedly coring a sample would move it to the left along its connectivity line in the figure, eventually causing it to stop percolating.

No simulations were run for $p < p_c$. Such a connection probability is unlikely for a natural material, because the rock-forming processes are affected by the percolation threshold. In a solidified foam such as tuff, the gas cannot escape until pores are sufficiently connected to percolate, but once some gas has escaped much pressure is relieve, and further connections may not be forced. Likewise, a crystalline rock under stress may crack until the fracture network reaches p_c ; at that point much stress is relieved, and further cracking is less likely. Because natural materials immediately below the percolation threshold are not commonly encountered, samples with width $>$ height are not generally needed; accordingly, simulations with width $>$ height were only rarely run.

The conclusion to be drawn from the work on this project is that, when working with low-connectivity materials, properties are best measured on samples for which width \approx height. Other shapes may bias the results.

Upscaling

This study did not focus on upscaling issues, but a simple comment is in order. When the correlation length is small – connectivity is high – percolation issues can largely be ignored.

This is because, at scales $L \gg \chi$, the material behaves classically. But when the correlation length is non-negligible, such that sample measurements are influenced by it, then extrapolating to larger sizes is non-trivial. A corollary is that when measuring diffusion (for example), long experimental times should be used, and the early-time data may be suspect.

Episodic flow

One episodic flow experiment was run, using a total of four episodes. Each episode consisted of a brief fracture-flow period, followed by a longer period during which the fracture was air-filled. Each episode introduced two new tracers: one sorbing and one non-sorbing. The reason for using so many tracers was simply that two kinds of analysis were used: effluent concentration measured during the experiment, and resident concentration measured destructively at the end of the experiment. Because simulations are not at risk from destructive sampling, only a single episode's worth of tracers was needed; additional episodes could then follow, with the effluent concentration and the tracer locations in the matrix recorded as desired.

A program for simulating tracer movement under episodic flow conditions was developed, but no simulations were run beyond a few small test suites. The CPU requirements exceeded the available capacity to obtain reasonably complete results. Nonetheless, the program development is discussed in the scientific notebook, and the source code and scientific notebook have been provided to LLNL separate from this report. It is recommended that episodic flow be considered for future work in this subject area.

Resulting Publications

Hu Q, RP Ewing, CI Steefel, L Tomutsa, and GB Hudson (2005). Multiple approaches to studying diffusion processes in geological media. *Geochim. et Cosmochim. Acta*, **69** A171, Suppl. 1.

Hu Q, RP Ewing, L Tomutsa, and MJ Singleton (2006). Pore connectivity, episodic flow, and unsaturated diffusion in fractured tuff. pp 70-76 *In Proc. 11th Internat'l High-Level Radioactive Waste Management Conference (IHLRWM)*, April 30 - May 4, 2006, Las Vegas, NV.

References

- Akaike H (1974). A new look at the statistical model identification. *IEEE Trans. Auto. Control* **19** 716–723.
- Bruce RR and A Klute (1956). The measurement of soil moisture diffusivity. *Soil Sci. Soc. Am. Proc.* **20** 458-462.
- Chow GC (1960). Tests of equality between sets of coefficients in two linear regressions. *Econometrica* **28** 591-605.
- Ewing RP and SC Gupta (1993). Modeling percolation properties of random media using a domain network. *Water Resour. Res.* **29** 3169-3178.
- Ewing RP and R Horton (2002). Diffusion in sparsely connected porespace: Temporal and spatial scaling. *Water Resour. Res.* **38** 10.1029/2002WR001412.
- Ewing RP and B Berkowitz (2001). Stochastic pore-scale growth models of DNAPL migration in porous media. *Adv. Water Resour.* **24** 309-323.
- Hunt, AG (2006). Scale-dependent hydraulic conductivity in anisotropic porous media from dimensional cross-over. *Hydrogeol. J.* **14** 499-507.
- Hunt, AG (2005). *Percolation Theory for Flow in Porous Media*. Lecture Notes in Physics 674. Springer-Verlag, Berlin.
- Julious SA (2001). Inference and estimation in a changepoint regression problem. *The Statistician* **50** 51-61.
- Meakin P and JM Deutch (1986). The formation of surfaces by diffusion limited annihilation. *J. Chem. Phys.* **85** 2320–2325.
- Philip JR (1957). The theory of infiltration 4: Sorptivity and algebraic infiltration equations. *Soil Sci.* **84** 257-264.
- Stauffer D and A Aharony (1994). *Introduction to Percolation Theory* (2nd ed.). Taylor and Francis Ltd., London.
- Zhou Q, HH Liu, GS Bodvarsson, and FJ Molz (2006). Evidence of multi-process matrix diffusion in a single fracture from a field tracer test. *Trans. Porous Media* **63** 473-487.

Appendix 1: Percolation Theory

“Percolation [theory] describes properties related to the connectivity of large numbers of objects which individually have some spatial extent, and for which their spatial relationships are relevant and statistically prescribed” (Hunt, 2005). More concretely, percolation theory describes the global properties arising in a system composed of many roughly equivalent constituent parts, where the main property that changes is the local degree of connection between those constituents. For example, suppose there is a large (preferably infinite) container filled with a random mixture of glass marbles and steel balls: what fraction of the spheres must be steel in order for an electrical current to pass through an arbitrarily large distance in the container? and, what is the system’s electrical conductivity as a function of the fraction of spheres that are steel?

Percolation theory works best near the transition between conducting and not-conducting; this is called the percolation threshold, or criticality. Below this point, there is no transport (or at least, no transport across a non-trivial distance); above that point, transport can occur. It turns out that just above criticality, many different systems display identical behavior. The various properties – mean size of a connected cluster, relative “mass” of the “infinite” cluster, conductivity, and so on – tend to follow power laws in the distance from criticality: for example, $\sigma(p) \sim (p - p_c)^2$, where σ is electrical conductivity, p is the proportion of the constituents that are active (e.g., the steel balls), and p_c is the critical proportion. In this equation, it is known that the conductivity exponent is 2.0 in 3D; in other dimensions it takes other values. This convenient (and mysterious) aspect of percolation theory – that equations change only in the value of their exponent when the dimension is changed – is called “universality”.

A fundamental variable in any system above but near the percolation threshold is the correlation length. Denoted χ , this distance scales as $\chi \sim (p - p_c)^{-\nu}$ (where $\nu = 0.88$ in 3D) (Fig. A1). A




Fig. A1 (from Ewing & Horton, 2003b). The correlation length χ is referred to here as the depth of the edge effect. The dashed line has an exponent of -0.88, the theoretical value.

system at criticality (i.e., at $p = p_c$) appears fractal, and fractals have no characteristic length. But above the percolation threshold, χ is the characteristic length. Processes operating at a scale smaller than χ “see” a fractal system, while processes at a greater scale “see” a uniform system. One therefore observes fundamentally different behavior at these different scales, with the transition occurring as a process’s scale exceeds the system’s characteristic scale χ . This is precisely what is seen in the imbibition experiments: the slope of the imbibition curve changes at a distance χ from the inlet face.

Some other fundamental percolation concepts and relationships also show up in the experimental and/or simulation results:

- (1) Near criticality, the accessible porosity ϕ_a decreases with distance from the inlet face as $\phi_a(l) \sim l^{-0.49}$ for $l < \chi$, then remains




Figure A2. Pu_{242} concentration as a function of distance from the inlet face. Ca is given for reference, to show the constant background of the rock.

constant for $l > \chi$. The value -0.49 is the quotient β/ν , where the values of β and ν in 3D are known from theory and verified by experiment. The change in ϕ_a with distance from the inlet face is evident in concentration versus distance plots obtained via LA/ICP-MS (Fig. A2); simulations (Fig. A3) showed essentially identical patterns. In Figure A3, the porosity that is accessible but not on the infinite cluster (the porosity between the two curves) is composed of those finite clusters which happen to intersect the inlet face.




Figure A3. Two kinds of porosity as a function of distance from the inlet face. The lattice measured 256^3 , and $p = 0.25$. The steep drop at the far end is caused by the outlet face, at which similar patterns hold. ("Accessible porosity" here denotes porosity accessible from the inlet face: "inlet-accessible porosity".)

- (2) The “constant” value of the accessible porosity ϕ_a , for distances $l > \chi$ (e.g., for distances 100..200 in Fig. A3), is itself a function of the proximity to the percolation threshold. Specifically, $\phi_a \sim (p - p_c)^{-\beta}$, where $\beta = 0.41$ in 3D.
- (3) Near criticality, the effective diffusion coefficient D decreases with distance from the inlet face as $D(l) \sim l^{-1.8}$ for $l < \chi$, then remains constant for $l > \chi$. Some consideration of this phenomenon was recently reported by Zhou et al. (2006), who reported that simulations of a long-term fracture-matrix tracer test only matched the data when a thin higher- D fracture surface layer was assumed. Similar findings were earlier reported by Narasimhan.
- (4) The imbibition transition is affected by anisotropy in the sample. Sample anisotropy means that the sample exceeds the correlation length in at least one dimension, but not in all dimensions. So (for example) a tall, thin sample of an isotropic medium may be anisotropic: it may be taller than one correlation length, but its radius may be less than χ (Hunt, 2006). This anisotropy may cause imbibition to change slope in a tall thin sample, although it showed no such transition over the same distance in a sample that was more evenly sized in all dimensions. This was seen in experiments and simulations both.
- (5) The shortest diffusion pathway between two points is greater than the straight-line distance; the ratio of the two values is generally called the tortuosity τ (though there are different ways of defining it; some use the inverse of the value here described, and some the square). This value is typically assumed to be scale-invariant for a given material. But as a system approaches criticality, the tortuosity increases with distance as $\tau(L) \sim L^{D_{\min}}$ for $L < \chi$, then remains constant for $L > \chi$. The exponent D_{\min} takes the value 1.34 in 3D.

Appendix 2: The Statistics of Breakpoints

Experimental data occasionally present as a “broken” line: threshold responses are a common example. If the conceptual model supports a two-line interpretation, the next issue is to identify the breakpoint x_b . If there is a need to statistically justify the presence of a breakpoint, an appropriate test must also be found. Because both experimental and simulated imbibition experiments often produced broken-line data, it was necessary to find a rigorous way both to identify the breakpoint, and to evaluate its statistical significance.

Several attempts were made to identify the x -value that gave the lowest total error sum of squares in a two-line fit (henceforth denoted MinSS). When this was not successful, a Hidden Markov model (HMM) was considered, which posits an underlying but hidden process that switches the observable system between two or more discrete states. Several HMM variants were implemented, and they worked quite well, far better than earlier attempts that fit two straight lines independently.

Subsequent conversations with a colleague who is a mathematical statistician (Dr. David Meek, USDA-ARS-NSTL) revealed several weaknesses in the approach, the main one being that two straight lines were being fitted independently:

$$\begin{aligned} y_1 &= a_1 + b_1 x && \text{for } x < x_b \\ y_2 &= a_2 + b_2 x && \text{for } x > x_b \end{aligned}$$

Dr. Meek recommended fitting the two lines simultaneously, requiring three fit parameters rather than four. The method used (Julious, 2001) requires that the two lines intersect at x_b , so a_1 can be obtained from the other three parameters:

$$a_1 = a_2 + x_b(b_2 - b_1)$$

Using this three-parameter fit, the MinSS method was almost always the best fitting method, usually even better than the HMMs. This is the method used in the analysis documented in this report.

Investigations also showed that, if the wrong breakpoint is identified, then the statistics used to decide between a one-line and a two-line interpretation of the data are more likely to support the two-line interpretation. In other words, to make the best possible decision between the one-line and the two-line interpretation, one must first find the best possible breakpoint, which requires that one (temporarily) assumes the two-line interpretation; only then can its validity be properly tested. This sequence is emphasized because some find it counter-intuitive.

Various statistics can be used to support the final decision. Several were reviewed, including the Akaike Information Criterion (AIC; Akaike, 1974) and a simple t-test of the difference between the two slopes. In the end, Chow’s F statistic (Chow, 1960) seems most appropriate, because not only is it consistent (like most of the other statistics examined), but it also allows the user to set a significance level. (In the AIC, for example, the “significance level” is essentially hardwired; the difference between the AIC and some others is simply the value of the built-in significance value.)

The breakpoint data presented in this report, particularly in the “sample shape” section, were obtained as follows:

1. First, the simulations were run, producing an output file (named <something>.out).
2. This file was then read by the breakpoint identification program (brkpt.pas), which identified breakpoints using the 3-parameter minSS method, and calculated Chow's F statistic for the two-line interpretation; results were written to a breakpoint file (<something>.brk).
3. This file was then read into Excel, and subjected to three screens: a two-line interpretation was accepted only if (a) it was significant according to Chow's F at $p < 0.05$, (b) the second slope was at least 0.4, and (c) the difference in slopes was at least 0.15. These last two threshold values are somewhat arbitrary, based on examination of many combinations.

When an expression like “two distinct slopes” is given in the report, it refers to cases which passed these three tests.

Numbers of interest from this analysis are (a) the fraction of realizations meeting these criteria, (b) the segment slope(s), and (c) the breakpoint x_b . For an inlet face of the lattice at $x = 0$, the breakpoint x_b corresponds to percolation theory's correlation length χ in lattice units.